

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
10 January 2002 (10.01.2002)

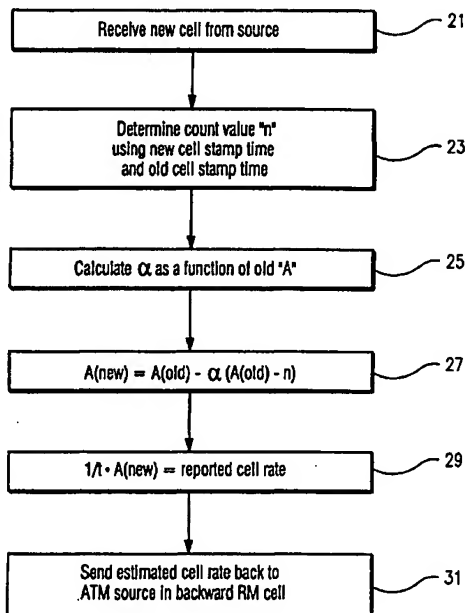
PCT

(10) International Publication Number
WO 02/03609 A2

- (51) International Patent Classification⁷: H04L 12/00
- (21) International Application Number: PCT/SE01/01510
- (22) International Filing Date: 29 June 2001 (29.06.2001)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
09/611,337 6 July 2000 (06.07.2000) US
- (71) Applicant: TELEFONAKTIEBOLAGET LM ERICSSON (publ) [SE/SE]; S-126 25 Stockholm (SE).
- (72) Inventor: WICKLUND, Rolf, Göran; Diligensvägen 82, S-131 48 Nacka (SE).
- (74) Agents: FORSSELL, Gunilla et al.; Albihns Stockholm AB, P.O. Box 5581, S-114 85 Stockholm (SE).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR ESTIMATING CELL RATE IN AN ATM NETWORK



(57) Abstract: A method/system for estimating cell rate is provided. The cell rate to be reported to an ATM (asynchronous transfer mode) source may be estimated using a previously calculated cell rate parameter weighted in a manner dependent upon whether a rate for the connection is high or low. An averaging parameter may be adjusted for each calculation so that for low rate connections newly measured data is more heavily weighted, and for high rate connections newly measured data is less heavily weighted. Relatively accurate and quick estimates for both high bandwidth and low bandwidth connections are achievable.



Published:

— without international search report and to be republished
upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

SYSTEM AND METHOD FOR ESTIMATING CELL RATE IN AN ATM NETWORK

This invention relates to a system and/or method for estimating a cell rate in an ATM (asynchronous transfer mode) network. More particularly, this invention relates to a system and/or method for estimating cell rate by exponentially averaging the
5 number of cell intervals between two cells received on a connection in an ATM network.

BACKGROUND

The increasing interest for high band services such as multimedia applications,
10 video on demand, video telephone, cellular phone services, and teleconferencing has motivated development of the Broadband Integrated Service Digital Network (B-ISDN). B-ISDN is based on ATM, and offers considerable extension of telecommunications capabilities.

ATM is a packet-oriented transfer mode which uses asynchronous time division
15 multiplexing techniques. ATM packets are called "cells" and have a fixed size. An ATM cell typically consists of fifty-three (53) octets, five (5) of which form a header and forty-eight (48) of which constitute a payload or information part of the cell. The header includes two quantities which are used to identify a connection in an ATM network over which the cell is to travel, particularly the VPI (virtual path identifier) and
20 VCI (virtual channel identifier).

At its termination or end points, an ATM network is connected to terminal equipment such as ATM network users (i.e. end systems). Between ATM end systems are a plurality of switching nodes having ports which are connected together by physical transmission paths or links. Thus, an ATM network is a collection of end systems
25 interconnected through one or more ATM switches. On a connection, an end system both sends data to other end system(s) on the network involved in the connection, and receives data from other end systems on the network involved in the connection. When an end system is a data sender, it is called a "source" and when a receiver it is called a

"destination." For example, see U.S. Patent Nos. 5,963,553 and 5,946,309, the disclosures of which are hereby incorporated herein by reference.

Multiple service categories of differing traffic characteristics and/or service guarantees have been defined for ATM networks. Examples include CBR (constant bit rate), rt-VBR (real-time variable bit rate), UBR (unspecified bit rate), nrt-VBR (non-real-time variable bit rate), and ABR (available bit rate).

ABR is intended primarily for non-real-time data applications with varying and/or unknown bandwidth requirements and which cannot easily be characterized in terms of a peak cell rate, a sustainable cell rate, and a maximum burst size. The source of an ABR VC (virtual circuit or connection) periodically creates and sends control cells called RM (resource management) cells which travel through the same path as data cells (e.g., traffic cells) of the VC to the destination of that VC. The destination then loops these cells back to the source through the same path. When an RM cell is traveling from the source to the destination it is called a forward RM cell; and when travelling from the destination to a source it is called a backward RM cell. Fig. 1 illustrates an exemplary RM cell including fields thereof.

Referring to Fig. 1, certain of these RM cell fields are read-only fields while others may be modified by intermediate ATM switches and/or the destination. For example, the CCR (current cell rate) field is a read-only field that contains the value of ACR (allowed cell rate) at the transmission time of the RM cell from the source. The DIR field represents a direction bit. When a source creates an RM cell, it sets DIR to zero (0) indicating a forward RM cell. Before the destination loops back an RM cell, it changes DIR to one (1) indicating a backward RM cell.

The ACR (contained in the CCR field) represents the rate the source is using to control its cell transmission for a VC or other network connection. The ACR value of a particular forward RM cell is inserted into that cell's CCR field at the time of its transmission. Sources adjust ACR based upon cell rate feedback carried by returning RM cells.

ABR switch mechanisms often suffer from CCR-reliance problems as follows. ER (explicit rate) values are determined by an ABR switch for each ABR connection. These values are sent back to the source in backward RM cells. ABR switch mechanisms typically read the CCR fields of RM cells passing through them and use these CCR fields in their ER calculations.

An implicit assumption that existing systems make is that the CCR field read by a switch from an RM cell on a given VC represents the actual rate of the VC at the moment of reading the field. This is often not true for at least the following reasons. First, CCR represents the maximum rate that the VC can use (i.e. ACR) at the time the RM cell was transmitted by the source. However, an ABR VC need not use all of its ACR. Thus, ACR may be larger than the actual rate of the VC. Second, when an RM cell is received by an intermediate switch and its CCR field read by the switch, the CCR field is old at that point and may not reflect the current value of ACR back at the source, not to mention the current actual cell rate. Third, a bad (or smart) source may insert a wrong ACR value into its CCR field in hope of acquiring more rate by doing so. This can cause the network to suffer in terms of buffer overflows; and fairness issues may arise.

To avoid the aforesaid CCR-reliance problems, per ABR VC rate measurements may be performed. One way of measuring (or estimating) ATM cell rate is to define a measuring time period "T" (e.g. 5 ms), and count the number of received ATM cells on a connection during that period. The cell rate is then n/T , where "n" is the count. This method may be rather accurate when the value of "n" is large (i.e. for high cell rate connections), because a minor deviation in the count number "n" does not substantially change the calculated rate. For example, if "n" changes from 75 to 73 when T is 5 ms, this only changes the measured cell rate from 15000 cells/s to 14600 cells/s, which is not a substantial change.

On the other hand, for low rate connections, the value "n" is small and even a single deviation in count value "n" can substantially affect the calculated cell rate. For example, assume the cell rate is just below one cell per period. In some intervals the

count "n" will be one (1) and the rate 200 cells per second, but in other intervals the count "n" will be zero (0) and the calculated rate will be zero (0). Likewise, even a shift in "n" from 1 to 2 will substantially change the measured cell rate from interval to interval. In short, this method of estimating cell rate is undesirable for low cell rate applications/connections.

Another manner in which to measure or estimate ATM cell rate is to measure the time (using a count of ATM cell intervals) between two cells received on the same connection (e.g. VC). The estimated cell rate in this case is $1/(n*t)$, where "t" is the length of a cell interval and "n" is the count value. This method has the opposite problem as that of the first measurement method previously described. This method is rather accurate for low cell rate connections ("n" is large), but does not work well for high cell rate connections ("n" is low). In high cell rate connections where "n" is low, a change in "n" from one (1) to two (2) would have a substantial effect on the measured cell rate. For example, if a rate over five or so cells received on a connection was truly about $1/1.1*t$ (but "n" varied between 1 and 2), the calculation of cell rate when "n" was 2 would be $1/2*t$ and would not accurately reflect the true cell rate.

The point where the two aforesaid methods have approximately the same accuracy is when "n" is equal in both methods, which is when "n" is equal to the square root of N (where N is the number of cell intervals in the measuring period).

One solution to the aforesaid problem(s) would be to use the first method (i.e. cell rate = n/T) for connections with high cell rates and the second method (i.e. cell rate = $1/(n*t)$) for low cell rate connections. The problem with this approach is that an ATM connection (e.g. VC) may have very different cell rates at different times, so both types of measurements may have to be performed on all connections at all times.

What is needed therefore, and an object of this invention, is an efficient method/system for estimating ATM cell rates in both high and low cell rate connections.

SUMMARY OF THE INVENTION

A method/system for estimating cell rate is provided. The cell rate to be reported to an ATM (asynchronous transfer mode) source may be estimated using a previously calculated cell rate parameter weighted in a manner dependent upon whether the previous rate was high or low. An averaging parameter may be adjusted for each calculation so that for low rate connections newly measured data is more heavily weighted, and for high rate connections newly measured data is less heavily weighted. Relatively accurate and quick estimates for both high bandwidth and low bandwidth connections are thus achievable.

For example, in one non-limiting embodiment the cell rate (CR) on an ATM connection may be estimated as $CR = 1/[A(\text{new}) * t]$, where: "t" is the length of a cell interval, and $A(\text{new}) = A(\text{old}) * (1 - \alpha) + (\alpha * n) = A(\text{old}) - \alpha * (A(\text{old}) - n)$. In the latter equation, "n" is a value indicative of the number of cell intervals between two received cells on the same connection; A(old) is the previously stored and calculated cell rate parameter A; and "α" is an averaging parameter to be adjusted between 0 and 1 as a function of cell rate on the connection. When "n" is large (low cell rate) on a connection, parameter "α" is adjusted to a high value(s) so that previously estimated or calculated cell rate parameters (e.g. A, "n", and/or CR) are soon forgotten and the newly estimated cell rate is based primarily upon newly measured data. However, when "n" is small (high cell rate) and may thus vary substantially relative to its previous value, parameter "α" is adjusted to a lower value so that newly measured data does not have a significant immediate effect on the estimated cell rate (CR) to be sent back to a source in a backward RM cell.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, features, and advantages of the invention will be apparent from the following more particular description of preferred embodiments as illustrated in the accompanying drawings in which like reference numerals indicate like

parts throughout the several views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating principles of the invention.

Fig. 1 shows various fields of an ATM RM cell.

Fig. 2 is a schematic illustration of a portion of an ATM network including first
5 and second ATM switches through which a plurality of sources and destinations communicate via ATM cells.

Fig. 3 is a schematic illustration of an ATM switching mechanism including a pair of ports for accommodating ATM cell flow.

Fig. 4 is a flowchart generally illustrating a manner in which new cell rate (CR)
10 is estimated upon receipt of an ATM cell according to an embodiment of this invention.

Fig. 5 is a flowchart illustrating how cell rate parameter A(new) is calculated in accordance with an embodiment of this invention.

Fig. 6 is a block diagram illustrating circuitry used for calculating A(new) in accordance with an embodiment of this invention.

15 DETAILED DESCRIPTION OF THE DRAWINGS

In the following description, for purposes of explanation and not limitation, specific details are set forth such as particular architecture, interfaces, circuitry, techniques, etc. in order to provide a thorough understanding of the present invention. However, it will be apparent to those skilled in the art that the present invention may be
20 practiced in other embodiments that depart from these specific details. In other instances, detailed descriptions of well known devices, circuits, graphs, and methods are omitted so as not to obscure the description of the present invention with unnecessary detail.

Fig. 2 is a schematic illustration of an ATM network including ATM sources A-
25 E which communicate with ATM destinations A-E via ATM switches SW0 and SW1. On a given ABR VC, a source sends data cells to a destination through one or more switches SW, with the source periodically creating and sending forward RM cells (see

Fig. 1). The aggregate traffic sent by the source is dynamically shaped to an ACR (allowed cell rate) which is controlled by the network. In particular, the source inserts its current ACR in the CCR field of forward RM cells. It also indicates its desired rate through the ER (explicit rate) field. The ER field is often set to the peak cell rate (PCR) of the connection. When a destination receives a forward RM cell, it changes the DIR bit to one (1), reduces the ER field to whatever value it can support, and loops the RM cell back to the source. The backward RM cells travel through the same path as that of the data and forward RM cells on the same VC, but in reverse direction.

In accordance with the ATM Forum Traffic Management Specification, intermediate switches SW have the option of intercepting forward and/or backward RM cells to reduce their ER fields to whatever values they can support so long as fairness is maintained. Thus, when an RM cell finally returns to its originating source, its ER field reflects the maximum possible rate on the most congested link along the connection's path. The source may then adjust its ACR using:

$$ACR_{new} = \min(ER \text{ in returning RM cell}, ACR_{current} + RIF * PCR)$$

Equation (1)

where RIF (rate increase factor) is a parameter that is determined during connection set-up. It can be seen that when $RIF=1$, ACR is totally controlled by the ER values in returning RM cells. Thus, ER values received by a source via returning RM cells have an impact on how the source sets its future cell rate parameters.

Certain embodiments of this invention relate to an ATM entity (e.g. ATM switch or end system) calculating or estimating an ATM cell rate that is to be used in calculating an ER value in a backward RM cell (e.g., measured CCR values need not be directly reported back to the source; but instead may be used as inputs to an algorithm that calculates the ER value, i.e., $ER = \text{function of measured CCR}$). By more efficiently calculating this value, the aforesaid CCR-reliance problems are addressed and/or resolved. ATM sources may then utilize these returned ER values in setting their ACR/ER values in the context of the ABR service category.

While the below-listed description and illustrations herein are in the context of an ATM ABR service category, it will be recognized by those of skill in the art that this invention is not so limited, and may be used or implemented whenever a measurement or estimation of the cell rate for a connection is needed (not just for ABR ER calculations).

As shown in Fig. 3, an ATM switch fabric connects a pair of ports 15, each having a queue 17 at its egress. The data cell and RM cell flows on an ABR VC are shown in differing solid lines with arrows. These lines may lead to any of the sources/destinations shown in Fig. 2. The method / system for calculating an ATM cell rate set forth herein may be carried out at an ATM switch, or alternatively at any destination in the network that receives ATM cells.

In certain embodiments of this invention, the rate calculation or estimation is based on an exponential averaging of the number "n" of cell intervals between two cells received on the connection (i.e. an improvement of the cell rate = $1/(n*t)$ methodology described previously). A cell rate (CR) estimate is calculated to be:

$$CR = 1/(A(new)*t) \quad (2)$$

where "t" is the length of a cell period and cell rate parameter A(new) is a function of "n" that is calculated as set forth below. A(new) in this particular embodiment is not an integer, but a real value having a fractional part.

Each time an ATM cell is received (either by a switch or end system which carries out this invention), it is timestamped. The difference between this timestamp and the previous timestamp of a cell from the same connection shows the number of cell intervals between two cells from that connection. The difference between the timestamps is thus "n".

When a new ATM cell is received on the connection at issue, new cell rate parameter A(new) is calculated in Equation (3) as follows:

$$A(new) = A(old)*(1-\alpha) + \alpha*n = A(old) - \alpha*(A(old)-n) \quad (3)$$

In equation (3): "*" means multiply; A(old) is indicative of and proportional to the mean value of the number of cell intervals "n" between two consecutively received cells recently received on the same connection (or alternatively may be the previously calculated and stored value of cell rate parameter A); "n" is the difference between
5 timestamp values (or any other number) that is indicative of the cell interval between the two cells received on the same connection; and " α " is an averaging parameter that is set (as will be described hereinbelow) as a function of the cell rate on the connection (e.g. α may be set as a function of A, CR, and/or "n" in different embodiments of this invention). As can be seen, A (old) is indicative of and proportional to a number of cell
10 intervals "n" between two previously received cells on the same connection.

Averaging parameter α , set as a function of the connection's cell rate at issue, is preferably adjusted to a value that is consistently greater than zero (0) and less than or equal to one (1). In other words, $0 < \alpha \leq 1$. When " α " is set to a relatively small value, the newly measured value "n" (which indicates the number of cell intervals between the
15 newly received ATM cell on the connection and the previously received ATM cell on that connection) will have little or small effect on A(new), i.e., the function has a long memory. However, when " α " is set to a relatively large value (e.g. from 0.75 to 1.0), the newly measured value "n" will have a larger effect on A(new), and the older values of "n" and A(old) are quickly forgotten. As can be seen from the equation for A(new)
20 above, if " α " is set to equal one (1), then A(new) will be calculated as being the newly measured value "n" for the received ATM cell (i.e., no memory at all). Thus, when the system is confident that the newly measured value "n" for a received cell on a particular connection is an accurate reflection of the actual cell rate (e.g. when "n" is very large), then " α " may be set to equal one (1).

25 For purposes of example only as to how "n" may be calculated, consider the following example. The term "n" is indicative of the number of received cells between two received cells belonging to the same connection. Suppose cells are received in the following order belonging to the number connection (each received cell includes a VPI/VCI value identifying its connection): 1, 2, 3, 4, 3, 2, 5, 6, 4, 5, 1. In this case, "n"

indicates that the number of cell intervals between the 1-cells is ten (10). Likewise, the number of cell intervals between the 2-cells is four (4).

For low cell rate connections, "n" is large and in and of itself is a good reflection of the cell rate. In such a case, averaging parameter " α " is set close to one (1), e.g. from about 0.5 to 1.0. Thus, the older A and "n" values will soon be forgotten and A(new) will be close or closer to the newly measured "n" value. However, for high cell rate connections, "n" may vary relative to itself in a significant manner between two intervals (e.g. "n" may be 1 for a given cell, and 2 for the next). Thus, in a high cell rate environment, " α " is set at a smaller level (e.g. from about 0.001 to 0.50) so that newly measured "n" value(s) do not have a significant immediate impact on A(new). In other words, for high cell rate connections, it is desirable to set " α " to a value so that "n" averaging takes place over a larger number of measured "n" values so that the reported rate is a more accurate representation of the true cell rate.

Accordingly, averaging parameter " α " is adjusted herein as a function of cell rate. Values "n", A and CR are reflections of cell rate. Thus, in certain embodiments " α " may be adjusted as a function of "n" while in other embodiments it may be adjusted as a function of A and/or CR for the connection at issue. If A (or "n" if it is the basis for adjustment) is small, then " α " is set to a relatively low value. If A (or "n") is large, then " α " is set to a relatively high value. In other words, averaging parameter " α " is adjusted so as to be proportional to cell rate (e.g. " α " is proportional to at least one of "n", A, and CR).

In certain non-limiting embodiments, averaging parameter " α " may be adjusted so as to be proportional to cell rate in accordance with the following equation:

$$\alpha = k * A(old) \quad (4)$$

where "k" is some constant and A(old) is some previously measured value of A (preferably the immediately previous value of A). In certain embodiments, " α " is set in a manner so as to be greater than some minimum value "m" such as zero (0), but no greater than one (1.0).

For very high cell rate connections, "n" as measured for a newly received ATM cell may vary between 1 and 2 for example. The more "1"s, the higher the true cell rate. Assume that a resolution of 0.1% is desired for the highest possible cell rate. When the count "n" is 2, the value A(new) (truly just above 1) will increase from one (1) to $(1 - \alpha) * A(\text{old}) + 2 * \alpha$ (or roughly $1 + \alpha$), which should not end up being higher than 1.001 in view of the desired 0.1% resolution requirement. A minimum value for averaging parameter " α " (for the highest cell rate) given the desired 0.1% resolution for the highest cell rate can thus be calculated as $1/1000$. This can be adjusted for implementation purposes to $1/1024$. Using such methodology, a table of values for parameter " α " can be developed in any suitable manner. An exemplary table in this regard is set forth below in Table 1, where " α " is set so as to be proportional to A(old).

TABLE 1: A(old) vs. " α "

	<u>A(old) value</u>	<u>"α" value</u>
	A < 2	1/1024
15	2-4	1/512
	4-8	1/256
	8-16	1/128
	16-32	1/64
	32-64	1/32
20	64-128	1/16
	128-256	1/8
	256-512	1/4
	512-1024	1/2
	A > 1024	1

In Table 1, values indicative of parameter " α " are a function of A for purposes of example only. A similar table may be used to set " α " as a function of a new or old "n" count value, or any other type of table may be used so long as " α " is adjusted as a function of cell rate or a cell rate parameter. As can be seen above in Table 1, the higher the cell rate, the smaller the value set for " α ." When a new ATM cell is received, after " α " has been calculated or adjusted in accordance with cell rate (as

exemplified above), the " α " value is inserted into equation (1) so that $A(\text{new})$ can be calculated and thereafter plugged into the equation for determining estimated cell rate CR.

As shown in Table 1 above, when A is greater than 1024 (low cell rate)
5 parameter " α " is set to one (1). This is because when "n" is large, it is a good reflection of cell rate when placed into the equation $1/(n \cdot t)$. Thus, when "n" or "A" exceeds a predetermined threshold then " α " may be set to a predetermined value such as $\frac{3}{4}$, 1, or the like (this value need not be 1 in all embodiments of this invention).

Fig. 4 briefly explains how cell rate estimate CR is calculated in one example
10 embodiment of the invention. Initially, a previous cell rate parameter $A(\text{old})$ and previous timestamp for the last cell received on a given connection are stored in memory. A new cell on the connection is received (step 21). Using the earlier timestamp and a new timestamp applied to the newly received cell, count value "n" is determined as described above (step 23). Averaging parameter " α " is adjusted or
15 calculated as a function of cell rate (e.g. as a function of $A(\text{old})$) (step 25). A new value $A(\text{new})$ is then calculated based in part upon effectively multiplying the newly determined "n" by averaging parameter " α ." For example, as shown in step 27, $A(\text{new}) = A(\text{old}) - \alpha (A(\text{old}) - n)$. After $A(\text{new})$ has been calculated, the estimated cell rate is then calculated: $CR = 1/(t \cdot A(\text{new}))$ (step 29). Following its calculation, the estimated
20 cell rate CR is used as input(s) to an algorithm that calculates the ER value, so that the source can efficiently adjust its ACR and/or ER values for future cells on the connection (step 31).

Cell rates in ABR ATM applications have been known to rapidly change. To cope with this potential, an extra check may be implemented to determine if a newly
25 measured "n" count value is of the same order of magnitude as the previously measured A value (i.e. $A(\text{old})$). If the new "n" value differs significantly (e.g. more than four (4) times A, or less than one-quarter of A), then parameter " α " should be set to a predetermined value (e.g. a mid-range value such as $\frac{1}{2}$) so that $A(\text{new})$, and thus CR, quickly converge to a new steady state value(s).

The above-listed equation (3) illustrates multiplication per se. From an implication perspective, it is not desirable to perform multiplication. However, effective multiplication by 2^k may be easily performed by logical shifts (i.e. shifts may be used to effectively multiply). Accordingly, in certain preferred embodiments of this invention, averaging parameter " α " may be set to the nearest value of 2^{-k} , $k=1, 2, 3, 4 \dots$. In such a manner, A(new) may then be calculated using subtractions and logical shifts as shown in Figs. 5-6. Thus, in accordance with the embodiment described above, averaging parameter " α " may be calculated as follows:

$$\alpha = (2^{**}[2\log A])/1024; \quad (5)$$

$$\alpha = \min(\alpha, 1); \quad (6)$$

$$\alpha = \max(\alpha, 1/1024); \quad (7)$$

$$\text{if } n > 4*A \text{ then } \alpha = 1/2; \text{ and} \quad (8)$$

$$\text{if } n < A/4 \text{ then } \alpha = 1/2. \quad (9)$$

In equation (5), "[2log A]" means the integer part of the 2-logarithm of A. In practice using the circuitry of Fig. 5 - Fig. 6, this means the position of the most significant binary bit of A. Thus, in equation (5), " $2^{**}[2\log A]$ " means in practice that all bits except the highest non-zero bit in A are cleared (note that A has a fractional part as well).

Given an assumption that 35 bits are available per connection for storage of values, the A value must have a good resolution and at the same time a wide range. Because 16 bits may not be enough to code it as an integer in all embodiments, it is preferably coded as a float (i.e. a float has a mantissa (m) and an exponent (e) as known in the art). For example, 9 bits may be used for the mantissa (m) and 5 bits for the exponent (e). The most significant bit of the mantissa (m) may be implicit, such that the value of A may be expressed as:

$$A = (1 + m/512) * 2^e \quad (10)$$

Equation (10) also is advantageous in that " α " is directly found as $\alpha = 2^{**e}/1024$. It is noted that before A is used in the calculation, it must be coded as a real number (20 bits integer part and 10 bits fractional part), and when the A(new) result has been calculated, it is again coded into float as shown in Fig. 5 - Fig. 6. Accordingly, there are 21 bits left over for the timestamp. However, one of these bits is preferably used to indicate "old" for a timestamp (see below), so 20 bits can thus be used for the actual timestamp for a received ATM cell. With 600 ns cell periods for example, this corresponds to 630 ms, or 610 bit/s. Timestamps may be maintained in number of cell intervals, i.e. if one cell interval is 600 ns, then a timestamp value of 10 corresponds to 6 microseconds.

It may happen that a source is silent for a period of time and thereafter begins again by sending a cell. In such a case, a cell received much earlier is preferably not be used due to its age. To handle this, all cells exceeding a predetermined age are identified by setting an "old" bit thereof for flagging purposes. For example, a background routine may check timestamps for all connections on a regular basis and marks those A(old) stamps that are too old (i.e. older than a predetermined value/age) with an old=1 stamp (i.e. set the old-bit in the same to "1"). In certain embodiments, an old timestamp (old=1) is applied to a cell who's stamp is more than 315 ms long (e.g. half the timestamp range). Whenever a cell is received, the old-bit is preferably reset. If the background routine is executed every 16th cell cycle, for example, and there are 32k connections, one loop with take 315 ms, which is what is required in certain embodiments.

Fig. 5 and Fig. 6 illustrate one manner in which the invention may be implemented, although it will be understood that this embodiment is not limiting as the invention may also be implemented in many other ways as will be appreciated by those skilled in the art.

Referring to Fig. 6, an ATM cell 49 is received by the ATM switch or end system. Also received is a channel number 50 which identifies the channel on which the cell was received. These signals 49 and 50 may come from another hardware (HW)

block that receives cells over a link and identifies the channels using the VPI/VCI values in the cells. New cell signal 49 clocks a running slot counter 52 into register 54. This is the new timestamp 56 for the channel on which the new cell was received. Channel number 50 is used to address memory 58 via controller 60, where data for each

5 channel (e.g. Old_time_stamp, mantissa (m) for A(old), exponent (e) for A(old), and old-bit indicative of whether data in RAM is too old for use) is stored. Controller 60 retrieves the most recently calculated and stored A(old) value (in float format) for the connection (e.g. VC) corresponding to the newly receive cell from memory 58 via its

10 (Old_time_stamp).

A background routine in controller 60 checks all old timestamp values of memory 58 versus a value from slot count 52 to make a determination as to whether they are too old to use. If so, the old bit of Old_time_stamp retrieved from memory 58 is set to 1 as discussed above (an indication that this stamp should not be used because

15 it is too old). Upon receiving a cell on a given connection, controller 60 initially checks to see if the old bit for the last timestamp for that connection is set to 1 (step 53). If so, the old-bit is reset to zero (0) (step 55) and A(old) is set to a very high value, preferably greater than 1024 by selector 62 and shift register 64 (step 55). In step 55, when A(old) is set to such high value, this results in a parameter α value of 1.0 (e.g. see Table 1

20 above) which in effect means that the A(old) value is not used in calculating A(new), i.e. CR is calculated solely based upon the newly measured count value "n" when placed in the equation $1/(n*t)$.

When it is determined in step 53 by controller 60 that the old bit is not set to 1, then A(old) is converted to a real number via shifting by register 64 and/or selector 62

25 (step 57). Controller then forwards the old timestamp value to comparator 66 which also receives from register 54 the new timestamp of the newly received cell. Comparator 66 determines the newly measured "n" count value in step 59 by subtracting the old timestamp from the new timestamp, and outputs the measured "n" value 68 to comparators 70 and 72. After "n" has been calculated by comparator 66,

controller 60 resets the "old timestamp" value to the "new timestamp" value which the controller received from register 54 (step 61). The new value for the "old timestamp" is then stored in memory 58.

Shift register 64 receives the mantissa and exponent bits of A(old) and converts
 5 the same to real, which is forwarded at 74 to comparators 70 and 72.
 Processor/comparator 72 receives "n" and real A(old), converts "n" to real in step 63, and then makes a determination as to whether "n" is greater than $4/A$ or less than $A/4$ in step 65 (to detect whether a rapid change in rate has occurred). If so (rapid rate change detected), then processor/comparator 72 causes k to be set to one ($k=1$) by selector 76
 10 (step 67). As discussed above, this results in " α " being set to approximately $\frac{1}{2}$ so that A(new) quickly converges to a new steady state estimate value. On the other hand, if it is determined that the step 65 query is negative, then k is set to the 2-logarithm of α (negated) in step 69 ($k = 10-e$) by arithmetic function 78. If k is then determined by function 78 at step 71 to be less than zero (0), then k is set to zero (0) at step 73 by
 15 function 78 and selector(s) 76, 80.

Thus, the output of selector 76 determines the k value sent to shift register 82. Register 82 also receives B from comparator 70, and performs step 75 by shifting B to the right by k bits. Thereafter, comparator 84 performs step 77 [$A(\text{new}) = A(\text{old}) - B$ (shifted)] to calculate A(new). Looking at step 75 followed by step 77, B has been set
 20 to $A-n$ and k has been set to $10-e$. Thus, A is from a practical standpoint set to: $A = A - (A-n) \gg (10-k)$ which is the same as $A(\text{new}) = (A(\text{old}) - (A^k-n)*2)/1024$, which is the same as the equation (3) listed previously.

After A(new) is calculated at 84, it is converted to float format at 86 in steps 79, 81 and 83 in a well known manner, and stored in memory 58 as A(old) replacing the
 25 previous value for the same.

After A(new) has been calculated, the estimated cell rate is then calculated: $CR = 1/(t*A(\text{new}))$. Following its calculation, the estimated cell rate CR is used to calculate ER (e.g., via a backward RM cell) so that the source can efficiently adjust its

ACR and/or ER values for future cells on the connection. Advantages include quick cell rate estimates for both high and low bandwidth conditions.

Fig. 6 is for purposes of example only and is not intended to be limiting. In other words, other types of circuits and/or software may be used to carry out the steps listed in Fig. 4 and Fig. 5.

Certain embodiments of this invention relate to an ATM entity (e.g. ATM switch or end system) calculating or estimating an ATM cell rate that is to be sent back to a source as an ER value in a backward RM cell. An averaging parameter may be adjusted for each calculation so that for low rate connections newly measured data is more heavily weighted, and for high rate connections newly measured data is less heavily weighted. Thus, estimates may be more efficiently calculated and relatively accurate and quick estimates for both high bandwidth and low bandwidth connections are achievable. By more efficiently calculating this value, the aforesaid CCR-reliance problems are addressed and/or resolved.

While the invention has been described in connection with what is presently considered to be the most practical and preferred embodiment, it is to be understood that the invention is not to be limited to the disclosed embodiment, but on the contrary, is intended to cover various modifications and equivalent arrangements included within the spirit and scope of the appended claims.

WHAT IS CLAIMED IS:

1. A method of calculating an estimated cell rate for a connection comprising the steps of:

receiving a cell on the connection;

5 adjusting an averaging parameter so that a value of the averaging parameter is a function of cell rate; and

estimating a cell rate for the connection using the value of the adjusted averaging parameter.

2. The method of claim 1, further comprising the step of determining a count
10 value indicative of a number of cell intervals between the received cell and a previously received cell on the same connection.

3. The method of claim 2, wherein said estimating step further includes estimating the cell rate for the connection using the count value indicative of the number of cell intervals between the received cell and the previously received cell on
15 the same connection.

4. The method of claim 1, wherein said adjusting step includes adjusting the averaging parameter so that the value of the adjusted averaging parameter is proportional to at least one of (a) a number of cell intervals between the received cell and a previously received cell, and (b) a previously stored value indicative of or
20 calculated at least in part based upon a number of cell intervals between two received cells.

5. The method of claim 4, further comprising the steps of: calculating a cell rate parameter $A(\text{new})$ in the following manner:

$$A(\text{new}) = A(\text{old}) * (1 - \alpha) + \alpha * n = A(\text{old}) - \alpha * (A(\text{old}) - n)$$

25 wherein α is the averaging parameter that is adjusted in said adjusting step, "n" is a count value indicative of a number of cell intervals between the received cell and a

previously received cell on the same connection, and A(old) is a previously calculated and stored value indicative of a number of cell intervals between two previously received cells;

- 5 estimating cell rate CR by $CR = 1/(A(\text{new}) * t)$, where "t" is a length of a cell period; and
 sending the estimated cell rate CR to a source via a backward RM cell.

6. The method of claim 1, further comprising the step of determining a count value indicative of a number of cell intervals between the received cell and a previously received cell on the same connection, and wherein said determining a count value step
10 includes comparing timestamp values for the received cell and the previously received cell.

7. The method of claim 1, wherein the received cell is an ATM cell, and the method further comprising the step of sending the estimated cell rate back to a source on an ATM network via a backward RM cell.

- 15 8. The method of claim 1, wherein said estimating step comprises adjusting the averaging parameter in a manner such that in a high cell rate connection a previously estimated cell rate parameter for the connection contributes more to the cell rate estimated in said estimating step than in a low cell rate connection.

9. The method of claim 1, wherein said adjusting step includes adjusting the
20 averaging parameter so that the parameter is larger for low cell rate connections than for high cell rate connections.

10. The method of claim 1, further comprising the steps of:
 determining whether a count value "n" relating to the received cell differs from a stored value by more than a predetermined amount; and when the count value "n"
25 differs from the stored value by more than the predetermined amount, then setting the averaging parameter to a predetermined value.

11. In an ATM network including one or more ATM switches, a method of estimating ATM cell rate on a given connection of the network comprising the steps of:

receiving an ATM cell on the connection;

adjusting an averaging parameter so that an adjusted value of the parameter is at

5 least partially a function of a value indicative of cell rate on the connection; and

estimating a cell rate for the connection using the adjusted value of the averaging parameter.

12. The method of claim 11, wherein said adjusting step includes adjusting the averaging parameter so that the adjusted value of the parameter is at least partially a

10 function of a previously stored value reflecting a number of cell intervals between two ATM cells received from a source on the connection.

13. An apparatus for estimating ATM cell rate for a connection, the apparatus comprising:

a receiver for receiving a cell on the connection;

15 a circuit for adjusting an averaging parameter to a value relating to cell rate of the connection; and

an estimating circuit for estimating a cell rate for the connection using the value of the averaging parameter.

14. The apparatus of claim 13, further comprising means for determining a
20 count value indicative of a number of cell intervals between the received cell and a previously received cell on the same connection.

15. The apparatus of claim 13, further comprising means for adjusting the averaging parameter so that the value of the averaging parameter is proportional to at least one of (a) a number of cell intervals between the received cell and a previously
25 received cell, and (b) a previously stored value indicative of or proportional to a number of cell intervals between two received cells.

16. The apparatus of claim 13, further comprising: means for calculating a cell rate parameter $A(\text{new})$ by:

$$A(\text{new}) = A(\text{old}) * (1 - \alpha) + \alpha * n = A(\text{old}) - \alpha * (A(\text{old}) - n)$$

wherein α is the averaging parameter, "n" is a count value indicative of a number
5 of cell intervals between the received cell and a previously received cell on the same connection, and $A(\text{old})$ is a previously calculated and stored value indicative of and proportional to a number of cell intervals between two previously received cells;

means for estimating cell rate CR by $CR = 1/(A(\text{new}) * t)$, where "t" is a length of a cell period; and

10 means for sending the estimated cell rate CR to a source via a backward RM cell.

17. In an ATM network including one or more ATM switches, an apparatus for estimating cell rate on a given connection of the network comprising:

means for receiving an ATM cell on the connection;

means for adjusting an averaging parameter so that an adjusted value of the
15 averaging parameter is larger for low cell rate connections than for high cell rate connections; and

means for estimating a cell rate for the connection using the adjusted value of the averaging parameter.

18. A method of calculating an estimated cell rate for a connection comprising
20 the steps of:

receiving an ATM cell on the connection;

adjusting an averaging parameter so that an adjusted value of the averaging parameter is larger for a low cell rate connection than for a high cell rate connection;
and

25 estimating a cell rate for the connection using the value of the adjusted averaging parameter.

1/5

RM (Resource Management) Cell Format [1]

FIELD	OCTET	BIT (s)	NAME	
Header	1-5	all	ATM Header	
ID	6	all	Protocol ID	
Type	DIR	7	8	Direction
	BN	7	7	BECN Cell
	CI	7	6	Congestion Indication
	NI	7	5	No Increase
	RA	7	4	Request/Acknowledge
		7	3-1	Reserved
ER	8-9	all	Explicit Cell Rate	
CCR	10-11	all	Current Cell Rate	
MCR	12-13	all	Minimum Cell Rate	
QL	14-17	all	Queue Length	
SN	18-21	all	Sequence Number	
Reserved	22-51	all	Reserved	
Reserved	52	16-11	Refer to I.371	
CRC-10	52-53	10-1	Refer to I.371	

Fig. 1




Fig. 2

2/5




Fig. 3

3/5

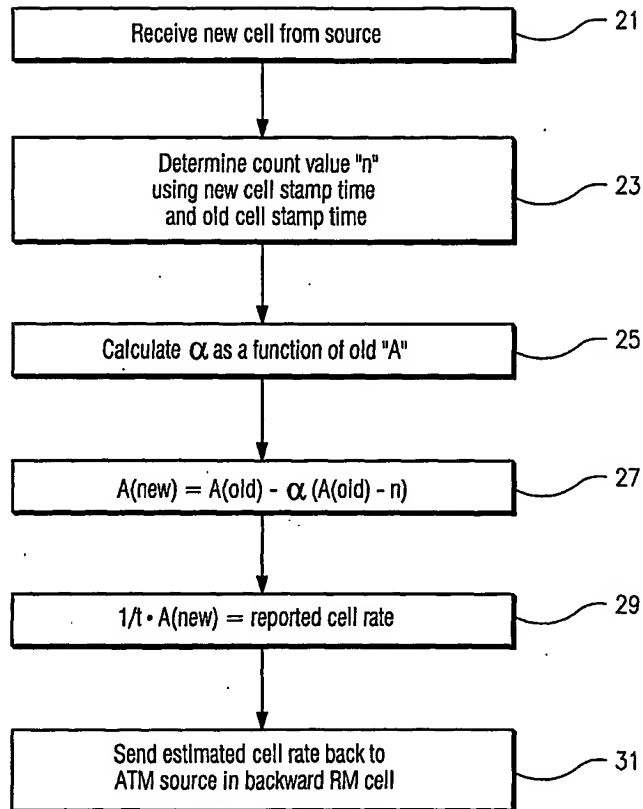


Fig. 4

4/5

Fig. 5

